Accepted for publication in the AIAA Journal of Guidance, Control, and Dynamics, 2004.

Lagrangian Delay Predictive Model for Sector-Based Air Traffic Flow*

Alexandre M. Bayen†, Pascal Grieder‡, George Meyer§, and Claire J. Tomlin¶

Abstract

We derive a control theoretic model of sector-based air traffic flow using hybrid automata theory. This model is Lagrangian, meaning that it models the properties of the system along its trajectories. A subset of this model is used to generate analytic predictions of air traffic congestion: we define and derive a *dynamic sector capacity* which we use to predict the time it takes to overload a given portion of airspace. This result links our approach with Eulerian models, which account for temporal variations of parameters in a fixed volume. We design and validate an air traffic flow simulator, to assess the accuracy of our predictions. The simulator is then used to show that flow scheduling and conflict resolution may be decorrelated under assumptions on aircraft density.

Keywords: scheduling, Air Traffic Control, network control, Lagrangian model, hybrid automata theory.

*Research supported by NASA under Grant NCC 2-5422, by ONR under MURI contract N00014-02-1-0720, by DARPA under the Software Enabled Control Program (AFRL contract F33615-99-C-3014) and by a Graduate Fellowship of the Délégation Générale pour l'Armement (France).

†Ph.D. Student, AIAA student member, corresponding author. Aeronautics and Astronautics, Durand 028, Hybrid Systems Laboratory, Stanford University, Stanford CA, 94305-4035. Tel: (650) 736-1423, Fax: (650) 723-3738, bayen@stanford.edu

‡Dipl. Ing. ETH, Ph.D. Student. Automatic Control Laboratory, Swiss Federal Institute of Technology, Physikstrasse 3, ETL K13.2, CH-8092 Zürich, Tel: +41 1 632-7313, Fax: +41 1 632-1211, grieder@aut.ee.ethz.ch

§Research Scientist, NASA Ames Research Center, MS 210-10, Moffett Field, CA 94035-1000, Tel: (650) 604-5750, Fax: (650) 604-0222, gmeyer@mail.arc.nasa.gov

¶Assistant Professor, Aeronautics and Astronautics, Courtesy Assistant Professor, Electrical Engineering. AIAA member. Hybrid Systems Laboratory, Stanford University, Stanford CA, 94305-4035. Tel: (650) 723-5164, Fax: (650) 723-3738, tomlin@stanford.edu

Report Documentation Page

Form Approved
OMB No. 0704-0188

1. REPORT DATE	2. REPORT TYPE	3. DATES COVERED
2004		**00-00-2004 to 00-00-2004**

4. TITLE AND SUBTITLE	5a. CONTRACT NUMBER
Lagrangian Delay Predictive Model for Sector-Based Air Traffic Flow	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Stanford University,Department of Aeronautics and Astronautics,Stanford,CA,94305**	8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)	10. SPONSOR/MONITOR'S ACRONYM(S)
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT
Approved for public release; distribution unlimited

13. SUPPLEMENTARY NOTES
AIAA Journal of Guidance, Control, and Dynamics, volume 28, number 5, pp. 1015-1026, Sept-Oct 2005. U.S. Government or Federal Rights License

14. ABSTRACT
see report

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**34**	

Symbol	Definition
a_i^0	initial arc length distance of aircraft i from the SFO airport along arrival route
$\vec{a}$	vector $[a_1^0, \cdots, a_N^0]$ of initial arc length distances for the N aircraft
b_i	variable used to compute the mode switching time of aircraft i
$\vec{b}$	vector $[b_1, \cdots, b_N]$ for the N aircraft
d_{min}^i	minimal distance from aircraft i to any other aircraft in the sector
dist_{LOS}	distance at which a LOS occurs
ΔL	requested outflow separation for merging traffic in the region of interest
ΔL_{in}	imposed inflow separation for cross traffic in the region of interest
ΔL_{out}	requested outflow separation for cross traffic in the region of interest
ΔT_{out}	outflow rate of the region of interest (one aircraft every ΔT_{out} time units)
ΔT_{act}	time period for one iteration of the simulator
ΔT_{LOS}^i	time until the next predicted occurrence of a loss of separation for aircraft i
$f(\cdot)$	penalty function for aircraft separation (associated to d_{min}^i for all i)
J	cost function encoding ATC controller action and sector state
$J_{\{\text{given action}\}}$	cost associated to a {given action} of the controller (VFS, shortcut, etc...)
M	Mach number
N_{max}	total number of aircraft in the sector of interest at a given time
n_{maneuver}	number of maneuvers the simulator can assign to an aircraft at any given time
N_{choice}	number of aircraft selected for analysis by the simulator
N_{moved}	total number of aircraft moved at a simulator iteration
n_{LOS}^i	number of LOS for aircraft i which would happen with a given set of maneuvers
N_{limit}	dynamic capacity of the sector of interest
R_ψ	rotation matrix of angle ψ for heading changes
T_{breach}^i	boundary condition breach time of aircraft i
t_i^{switch}	time at which aircraft i undergoes a mode switch by ATC
t_{block}	time at which a metering condition is imposed
T_{limit}	saturation time of the sector of interest
TOA_{pred}^i	predicted time of arrival of aircraft i (at TRACON)
TOA_{real}^i	actual time of arrival of aircraft i (at TRACON)
$\vec{v}_{\text{current heading}}$	velocity vector of a given aircraft at its current heading
v_{min}	minimum aircraft speed
v_{nom}	nominal aircraft speed
v_{max}	maximum aircraft speed
$w_{\text{given action}}$	weight (penalty) associated to a given action of the controller
$\vec{x}_i$	position of aircraft i
x_{ex}	end of the controlled area
x	distance to the destination airport
x_i^{switch}	location at which aircraft i undergoes a mode switch by ATC

Acronym	Definition
ARTCC	Air Route Traffic Control Center
ATC	Air Traffic Control
ATM	Air Traffic Management
ATCSCC	Air Traffic Control System Command Center
BC	Boundary Conditions
CONUS	Continental US
ETMS	Enhanced Traffic Management System
FACET	Future ATM Concepts Evaluation Tool
LWR	Lighthill-Whitham-Richard
LOS	Loss Of Separation
NAS	National Airspace System
nm	Nautical Mile
SFO	San Francisco Airport
TFM	Traffic Flow Management
TMU	Traffic Management Unit
TRACON	Terminal Radar Approach CONtrol
VFS	Vector For Spacing

1 Introduction

The *National Airspace System* (NAS) is a large scale, layered, nonlinear dynamic system: its control authority is currently organized hierarchically with a single *Air Traffic Control System Command Center* (ATCSCC), in Herndon VA, supervising the overall traffic flow. This is supported by 22 (20 in the continental US or CONUS) *Air Route Traffic Control Centers* (ARTCCs, or simply, Centers) organized by geographical region up to 60,000 feet [1, 2, 3, 4, 5, 6]. Each Center is sub-divided into about 20 sectors, with at least one air traffic controller responsible for each sector. Each sector controller may talk to 25-30 aircraft at a given time (the maximum allowed number of aircraft per sector depends on the sector itself). The controller is in charge of preventing *losses of separation* between aircraft, keeping them separated by more than 5 nautical miles horizontally, and 2000 feet vertically. In general, the controller has access to the aircraft's flight plan and may revise the altitude and provide temporary heading assignments, amend the route, speed, or profile, in order to attempt to optimize the flow and to keep aircraft *separated*, as well as to provide weather reports and winds. An illustration of the current control structure is presented in Figure 1.

Existing NAS modeling tools have functionality which spans the modeling of runway and airport capacity and operations, through airspace operations and conflict resolution [7, 8], to human factors and man-machine integration. See [9, 10] for detailed surveys of NAS modeling and conflict detection and resolution methods. A recent tool, FACET [11, 12], represents the first accurate NAS simulation tool, with the additional capability of a "playback" mode using actual traffic flow data. Our goal is to develop a model which complements existing tools, by providing a control theoretic component which models the influence of Air Traffic Control

(ATC). While the additional logic required to model the actions of the controller does not pose a significant computational problem if the aircraft density in the airspace is low, it becomes an issue as the density increases. The long term goal of increasing capacity as well as safety in the NAS cannot be achieved without an in-depth analysis of the applied control logic. If such a system were shown to model the current airspace with sufficient accuracy, then a wide array of applications would become feasible, including providing additional support for ATC in predicting delay.

In this paper we present a model, as well as analytic and simulation results, of the aircraft and controller actions within a sector of airspace. Our approach is *Lagrangian*, meaning that it accounts for the trajectories of the aircraft and parameters transported along them, for example, local aircraft density (average number of aircraft per surface unit for a given portion of airspace), momentum or average speed. Most of the existing NAS models use this framework, see for example [13, 14, 15, 11, 16, 17, 18]. However, we bridge our results with *Eulerian* approaches such as [19]. Eulerian approaches are control volume based: they account for temporal fluctuations of quantities in a given volume, for example, the number of aircraft present in a portion of airspace. The connection between these two approaches is made through a new concept which we introduce: sector *dynamic capacity*, also linked to other Lagrangian studies [14] and complexity metrics [20, 21]. We use our model to study the effect of aircraft flow density requirements at sector boundaries, due, for example, to miles-in-trail requirements at airports in subsequent sectors: the model allows for prediction capability as to how the current system might react to imposed flow conditions. In addition, the model allows for the testing of the effectiveness of different controller policies in minimizing delays in posted flight plans.

This article has two components: airspace modeling and analysis, and validation and simulation. In the first, we present a mathematical model for a controlled sector, based on a hybrid system model for each aircraft which encodes simple aircraft dynamics under the discrete action of the controller. We have observed that the set of commands used by controllers, while large, is finite, and consists of simple actions such as: "turn to heading of x degrees", "hold current heading", "fly direct to jetway y", "increase speed to z knots". We then analyze this model and define the concept of sector dynamic capacity, which we combine with analysis to predict the time it takes to overload given sectors of airspace, and thus predict delays, assuming controllers use a subset of their available control actions. We finally link these results with Eulerian approaches. In the second part of this article, we validate the previous results against real data. Since our results cannot be tested on the real ATC system directly, we design a simulator of the system, which we implement in C++ interfaced with MATLAB. This simulator consists of the mathematical model derived in the first part of the paper, augmented with a logic for switching between the different controller states. This logic models the actions of a human air traffic controller, and results from minimizing a particular cost function, which we derive. We validate this simulator against Enhanced Traffic Management System data (ETMS) and then show that our analytical predictions for sector capacity are effectively observed in simulations. Finally, we use our simulator to identify flow conditions under which conflict resolution decorrelates from *metering* problems (i.e.

scheduling aircraft according prescribed rules). This result has implications for numerous ATC flow management studies which rely implicitly on this assumption [22, 23, 19, 14].

The data that we present in this paper pertain to several sectors within the Oakland ARTCC, located in Fremont, CA. Our methodology, however, is general and would apply to any other En-Route portion of the NAS. The geographical data presented is available in the public domain (we used JEPPESEN [24] high altitude en route charts). Our controller model and cost function have been designed based on our observations, over several hours, of sector controllers for given sectors; we have made justified approximations where appropriate for the study in this paper. While our model is general, most of the scenarios considered in this paper do not represent normal traffic flow: because we are interested in modeling delay propagation of the system under stress, the traffic scenarios modeled represent heavy traffic flow along jetways.

The contributions of this paper are a new mathematical model for airspace sectors, which uses hybrid system theory, an analytical solution to the Lagrangian problem of delay propagation in the network, and its link with Eulerian approaches, and the concept and use of sector dynamic capacity. On the application side, the novelty is in the validation of a control theoretic model of the human air traffic controller, as well as in the validation of the analytical predictions against real data. Finally, the decorrelation results shown by our simulations are new.

This paper is organized as follows: Section 2 presents the model used for aircraft dynamics and air traffic controller actuation (Section 2.1). This model is used in Section 2.2 to predict the propagation of airspace congestion, and to define sector capacity. Section 3 presents the design of the simulator (Section 3.1), its use in validating our analytical predictions (Section 3.3), and its use in demonstrating the decorrelation between conflict resolution and flow metering (Section 3.4).

2 Air Traffic Flow modeling and analysis

The structure of the NAS is complex, for it involves a multitude of interacting agents and technologies: aircraft monitoring, flow management, communication, and human-in-the-loop. For the present work, in which we are interested in predicting delay, we extract and model only the features which are important for this purpose. We model a portion of the Oakland ARTCC, which contains five *sectors*. These sectors surround the Oakland TRACON (Terminal Radar Approach CONtrol), which controls the aircraft on their approaches into San Francisco, San Jose and Oakland airports. The TRACON is the final destination of the traffic that we consider in this paper.

We model a sector by a portion of airspace containing aircraft under the local control of the responsible air traffic controller (Figure 2). The interior of this domain is the controlled area (in which the local controller can actuate the flow). Within each sector, navigation infrastructure, including *jetways, waypoints* and *navigation aids*, is used to help the flow follow desired patterns; we therefore include it in the model and use it, even if it is observed that more than 40% of the aircraft may fly off the jetway at any given time. Our model allows

for aircraft to fly at different altitudes, but not to climb or descend. Altitude changes are not crucial for the effects we want to identify: the type of sector overloads we are interested in mainly result from aircraft acceptance rates at destination airports.

2.1 Aircraft behavior

We use a *hybrid model* for each aircraft. A hybrid model describes the evolution of a system by a set of *discrete modes*, each associated with a *continuous dynamical system*, and *discrete switches* which enable the system to jump from one mode to another instantaneously. In mathematical terms, we will describe the motion of aircraft i by:

$$\dot{\vec{x}}_i = \frac{d\vec{x}_i}{dt} = \vec{v}_{\text{current heading}} \tag{1}$$

where $\vec{v}_{\text{current heading}} \in \mathbf{R}^2$ is a constant velocity vector held by the aircraft until the next discrete switch: a heading or speed change which changes $\vec{v}_{\text{current heading}}$. $\vec{x}_i \in \mathbf{R}^2$ is the planar position of aircraft i. Integration of equation (1) over time produces a continuous piecewise affine trajectory. We prefer such a model over a continuous dynamical model for two reasons. First, the time scale of a "change in aircraft behavior" (for example a turn or slow down) is on the order of 30 seconds, whereas the time scale of a straight line portion of the flight is usually much longer, sometimes half an hour or more, thus we ignore the dynamics of such maneuvers and focus on their effects only (the set of resulting straight lines). Second, the update rate of ATC monitoring is in general not more than 30 seconds, which makes the details of these maneuvers inaccessible to the ATC. This approximation is widely accepted in literature [25, 26, 27, 17, 28, 22].

Monitoring ATC shows that a finite set of maneuvers, which depends on local parameters, is used. Combinations of these maneuvers result in a conflict-free flight environment in which the constraints of the air traffic flow are met. The maneuvers shown in Figure 4 are modeled (and consist of changing the right hand side of (1) according to certain rules which we now make explicit). The validity of models similar to this has been confirmed by statistical studies realized at MIT ICAT (Histon and Hansman [20]).

1. *Speed change:* ATC may decelerate or accelerate the aircraft along its flight plan:

$$\vec{v}_{\text{modified speed}} := \lambda \cdot \vec{v}_{\text{current heading}} \tag{2}$$

   where $\lambda \in \mathbf{R}^+$ defines the magnitude of the velocity change. Our model will allow a finite set of speeds (which means λ has a finite number of acceptable values). This encodes the fact that the ATC generally has a finite set of possibilities in the choice of speeds, because the aircraft flies at its optimal speed per altitude and ATC will speed up or slow down the aircraft by not more than 10% of the current value.

2. *Vector-for-spacing (VFS):* This maneuver consists of a deviation of the aircraft from its original flight plan for a short time (part 1 of the maneuver), and a second deviation,

bringing it back to its original flight plan (part 2 of the maneuver). This stretches the path that the aircraft must follow, and therefore generates a delay. The length of this maneuver depends on the geometry of the sector. Calling R_ψ the rotation matrix by angle ψ, we have:

$$\begin{array}{ll} \vec{v}_{\text{part 1}} := R_\psi \cdot \vec{v}_{\text{current heading}} & \text{(First half of the maneuver)} \\ \vec{v}_{\text{part 2}} := R_{-2\psi} \cdot \vec{v}_{\text{part 1}} & \text{(Second half of the maneuver)} \end{array} \tag{3}$$

3. *Shortcut / Detour:* In certain situations, the ATC will have the aircraft "cut" between two jetways, a maneuver which could either shorten or lengthen the flight plan. The decision to command such a maneuver is often dictated by conflict resolution, but could also be used to shorten the overall flight time if sector occupancy allows it (sometimes called "direct-to" by pilots):

$$\vec{v}_{\text{shortcut}} := R_\psi \cdot \vec{v}_{\text{current heading}} \quad \text{for the duration of the maneuver} \tag{4}$$

   until the next ATC action is taken. Here again ψ is the angle by which ATC turns the aircraft to achieve the shortcut.

4. *Holding pattern:* In some extreme conditions, holding patterns are used to maintain an aircraft in a given region of space before eventually letting it follow its original flight plan. This is modeled by assigning the aircraft to a predefined zone and keeping it there while preventing other aircraft from entering that zone.

2.2 Lagrangian analysis of delay propagation in the NAS

A large proportion of air traffic jams, i.e. portions of airspace saturated by aircraft, is generated by restrictions imposed at destination airports, usually themselves driven by weather or airport congestion. These restrictions are imposed as either miles-in-trail or minutes-in-trail, representing the distance (or time) required between aircraft in a flow incoming to the TRACON, and are referred to as *metering constraints*. Figure 5 illustrates the topology of the flows incoming to San Francisco Airport (SFO) which are often subject to this type of constraint. These constraints tend to propagate backwards from the airport into the network, and result in miles-in-trail constraints imposed at the entry points of each sector. For example, in the case of Figure 5, typically, the backpropagation of these metering conditions is as follows: TRACON $\rightarrow$ sector 34 $\rightarrow$ sector 33 $\rightarrow$ Salt Lake Center $\cdots$ and similarly for the two other flows.

In the current system, these restrictions are imposed empirically. We would like to understand *(i)* how the traffic jams propagate; and *(ii)* what the optimal control policy should be under these restrictions, in order to ensure maximal throughput into the TRACON.

2.2.1 Shock wave propagation

We present a simple Lagrangian model of merging flows, introduced earlier in [29, 30]. This model of merging flows predicts the backpropagation of a traffic jam from a destination

airport into the network. It is Lagrangian, because it models the trajectories of all agents in the considered space, rather than averaged quantities in a control volume, such as the number of aircraft in a given sector. This model can be applied to the merging flows of the type shown in Figure 5. We will use it here to derive the dynamic capacity of a sector. Given prescribed inflow and outflow conditions, we define the *dynamic capacity* of a sector to be the number of aircraft which can be actuated[1] in this sector, until ATC action has to occur upstream from this sector in order to not violate the outflow conditions. We assume in this definition that the sector is initially empty, but we will show how the definition may be applied to cases in which the sector is initially non-empty. The dynamic capacity is a concept which appears naturally in the following problem:

Given a required spacing between the aircraft (metering constraint) of ΔT_{out}, compute a controller policy which will force groups of aircraft to exactly *satisfy the metering constraint at the sector exit point (each aircraft separated by exactly ΔT_{out}) while maintaining separation at all times.*

We first consider a very simple version of the problem, in which the controller uses only two modes (fast and slow). This model is not overly restrictive: we can use several methodologies to map the full automaton of Figure 4 to this model: see for example [23, 14].

We introduce the initial arc length distance $a_i^0 \in \mathbf{R}$ of aircraft i along its arrival route to the airport. For example, $a_i^0 = -200$ means that aircraft i has to fly 200 nm before landing in the airport. We call $x_{\text{ex}} \in \mathbf{R}$ the location at which the metering condition is imposed. For example, $x_{\text{ex}} = -50$ means that the metering is applied 50 nm from the airport. It is possible to assume without loss of generality that the aircraft are numbered in order of arrival (the a_i^0 are indexed in increasing order of magnitude).

We consider the following situation: all aircraft are initially at maximum speed $v_{\max}$, and in order to enforce metering, ATC slows down aircraft i to its minimum speed $v_{\min}$ (see Figure 6), at a location x_i^{switch} and time t_i^{switch} which we will determine (see Figure 7). This scenario is represented as a dash-dot line in Figure 4. We impose the condition that each aircraft cross the metering point x_{ex} at exactly $t_{\text{block}} + (i-1)\Delta T_{\text{out}}$ where t_{block} is the time at which the metering condition is initiated. For simplicity, we take the origin of time at $t_0 = 0$. This leads to the following kinematic equations of the aircraft under this actuation:

$$\begin{aligned} x_i(t) &= a_i^0 + v_{\max} t \quad t \in [0, t_i^{\text{switch}}] \\ x_i(t) &= b_i + v_{\min} t \quad t \in [t_i^{\text{switch}}, t_{\text{block}} + (i-1)\Delta T_{\text{out}}] \end{aligned}$$

The assumption of continuity of $x_i(t)$ enables us to solve for b_i, from which we deduce the switching time as: $t_i^{\text{switch}} = (b_i - a_i^0)/(v_{\max} - v_{\min})$. Under the following feasibility conditions:

$$a_i^0 \in [x_{\text{ex}} - v_{\max}(t_{\text{block}} - (i-1)\Delta T_{\text{out}}), x_{\text{ex}} - v_{\min}(t_{\text{block}} - (i-1)\Delta T_{\text{out}})] \tag{5}$$

the propagation speed of the traffic jam may be computed analytically, by solving for the

[1] ATC *actuation* of an aircraft means any alteration of its original flight plan, in order to meet desired conditions in the region of interest.

location of the edge of the traffic jam in space and time:

$$\begin{array}{rcl} t_i^{\text{switch}} = & \frac{x_{\text{ex}} - v_{\min} t_{\text{block}} - (i-1)\Delta L - a_i^0}{v_{\max} - v_{\min}} \\ x_i^{\text{switch}} = & a_i^0 + \frac{v_{\max}[x_{\text{ex}} - v_{\min} t_{\text{block}} - (i-1)\Delta L - a_i^0]}{v_{\max} - v_{\min}} \end{array} \tag{6}$$

where $\Delta L := v_{\min}\Delta T_{\text{out}}$ is the metered spacing at the outflow of the sector. At a given time t, we call *traffic jam* or *metered platoon* the set of aircraft such that $t_i^{\text{switch}} \leq t$. These aircraft have already been actuated (see Figure 6). It follows directly from (6) that the traffic jam will not grow if the two following conditions are met:

$$\begin{array}{rcl} t_i^{\text{switch}} < t_{i+1}^{\text{switch}} & \Leftrightarrow & \Delta L < a_i^0 - a_{i+1}^0 \\ x_{i+1}^{\text{switch}} < x_i^{\text{switch}} & \Leftrightarrow & \left(\frac{v_{\min}}{\Delta L}\right) < \left(\frac{v_{\max}}{a_i^0 - a_{i+1}^0}\right) \end{array} \tag{7}$$

Condition (6) above is a sufficient condition for the length of the traffic jam to decay (which can be observed by inspection of the slope of the switching curve or *shock wave* of points $(x_i^{\text{switch}}, t_i^{\text{switch}})$ displayed in Figures 7 and 8). It is a local property of the problem: the conditions depend only on $a_i^0 - a_{i+1}^0$, not on all aircraft considered here. The second equation in (6) is in fact a one-dimensional discretized steady Lighthill-Whitham-Richard (LWR) equation, which appears naturally in highway congestion problems [31]. This fact links this Lagrangian approach, based on aircraft trajectory analysis, with Eulerian approaches such as [19], which are based on conservation equations; it relates local properties of the flow (local monotonicity of a variable, here the x-location of the wavefront) to global quantities (here the trajectories of the aircraft). This result is illustrated in Figure 7.

The condition that each aircraft reaches x_{ex} exactly at the time prescribed is restrictive, what is important is the flow rate but not the actual crossing times. Therefore it would be of greater use to pose the problem as follows: *Given* $\{a_i^0\}_{i\in[1,N]}$*, compute the switching policy which delivers at most one aircraft every* ΔT_{out} sec. *at* the location x_{ex} *while maintaining separation, and minimizes the arrival time of aircraft* N. This problem may be posed as a linear program: (**a.**) minimize the arrival time of aircraft N, (**b.**) while separating the aircraft by more than ΔT_{out} at x_{ex}, (**c.**) with at most one switch between the initial position $a_i^0 \leq x_{\text{ex}}$ and the exit x_{ex} of the considered airspace:

a. Minimize $\quad [0, \cdots 0, -1]\vec{b}$

b. Subject to

$$\begin{bmatrix} -1 & 1 & 0 & \cdots & \cdots & 0 \\ 0 & -1 & 1 & \ddots & & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & & \ddots & -1 & 1 & 0 \\ 0 & \cdots & \cdots & 0 & -1 & 1 \end{bmatrix} \vec{b} \succeq v_{\min} \begin{bmatrix} \Delta T_{\text{out}} \\ \vdots \\ \vdots \\ \vdots \\ \Delta T_{\text{out}} \end{bmatrix}$$

c. $\quad \vec{a} \preceq \vec{b} \preceq \frac{v_{\min}}{v_{\max}} \overrightarrow{a} + \left(1 - \frac{v_{\min}}{v_{\max}}\right) x_{\text{ex}}[1, \cdots, 1]^T$

where $\vec{a} = [a_1^0, \cdots, a_N^0]^T$ and $\vec{b} = [b_1, \cdots, b_N]^T$. Note that the right hand side of (**b.**) can be changed to $[\Delta T_1, \cdots, \Delta T_N]^T$ in order to account for time-varying metering conditions. The

advantages of this formulation are that one includes the possibility to optimize an objective function, which is in this case the arrival time of the last aircraft in the platoon, and that one is able to deal with time varying acceptance rates at the airport. Conditions (7) for shock monotonicity derived previously are still valid locally for any solution derived with the linear program above.

2.2.2 Sector overload predictions

Using the analysis of the previous section, it is fairly easy to predict the *dynamic capacity* of a sector. Consider the worst case scenario: an incoming flow of aircraft, each at $v_{\max}$, separated in time by ΔT_{in} chosen to violate the second condition in (7). This will create a traffic jam originating at SFO (equated with entrance to SFO TRACON), which "piles up" and progressively fills sector 34. We call l the arc length distance of the portion of the arrival jetway contained in sector 34 (see Figure 6). We assume that the sector is initially empty. Using equations (6), we can compute the maximum number N_{limit} of aircraft which can be stacked along the length l of the jetway in the sector, before no space is available anymore on this jetway. These aircraft are labeled as *metered platoon* in Figure 6; for example, for the situation shown in this Figure, approximately half of l is occupied by the metered platoon at the time considered, so the number of metered aircraft is approximately $N_{\text{limit}}/2$. When the number of metered aircraft reaches N_{limit}, after a time called T_{limit}, the actuation shown in Figure 6 (ATC slows aircraft down) has to occur upstream (in sector 33). N_{limit} and T_{limit} are respectively given by:

$$N_{\text{limit}} = \frac{l(v_{\max} - v_{\min})}{v_{\max} v_{\min}(\Delta T_{\text{out}} - \Delta T_{\text{in}})} \tag{8}$$

$$T_{\text{limit}} = \frac{l}{v_{\max} v_{\min}} \frac{v_{\max}\Delta T_{\text{in}} - v_{\min}\Delta T_{\text{out}}}{\Delta T_{\text{out}} - \Delta T_{\text{in}}} \tag{9}$$

If aircraft are initially present in the sector, these two quantities can be modified by replacing l by the distance of the last aircraft arrived in the sector to the entrance of the sector. We call N_{limit} *dynamic* capacity (and not static) because it depends on the inflow and outflow conditions (and not only on geometric parameters).

Several comments can be made regarding the two previous results: *(i)* As $v_{\min} - v_{\max} \to 0$, $N_{\text{limit}} \to 0$: no aircraft can be handled in the sector because no actuation is possible (it is not possible to "make the aircraft lose time" in this sector); *(ii)* As $\Delta T_{\text{out}} - \Delta T_{\text{in}} \to 0$, $N_{\text{limit}} \to \infty$ and $T_{\text{limit}} \to \infty$: if the incoming flow is such that it is almost metered as imposed at the exit of the sector, the number of aircraft required to saturate this airspace becomes large and the time it takes to saturate this sector grows accordingly.

The construction of the switching curve $(x_i^{\text{switch}}, t_i^{\text{switch}})$ described previously, can be used to compute the maximal extent of a traffic jam, or the portion of jetway affected by a traffic jam. Using (6), one can trace the shock location in the (x, t) plane (Figure 8). The edge of the traffic jam, called x_m, obtained at t_m gives the worst situation obtained from the initial

configuration a_i^0 of the aircraft: at t_m, the extent of the traffic jam is at its maximum. In the case of Figure 8, we see that the traffic jam does not propagate more than 300 nm upstream from the destination of the aircraft, called x_{ex}. Therefore no metering conditions need to be applied upstream from that point. In the current system, such information is not available to the ATC, thus leading to extra buffers taken by the controllers, which in turn leads to non optimal operating conditions as well as backpropagation of "virtual overload", a set of conservative precautions.

3 Validation against ETMS data

The models of the previous section rely on a mathematical analysis of metering, based on geometry. In order to understand how realistic these models are, we need to validate them against the real system. Since the real system is not available as a testbed, we have designed an abstraction of it which reproduces its behavior adequately. In this section, we present the design of a simulator which mimics true ATC behavior, and the validation of this simulator against ETMS data. The graphhical interface of the simulator is shown in Figure 3. In Section 3.3, we validate the analytical predictions of Section 2, using the simulator, and in Section 3.4, we show how the simulator may be used to derive conditions on the decorrelation of flow metering and conflict resolution.

3.1 Simulator design

Our simulator is based on empirical studies that we realized at the Oakland ARTCC, its core is based on observed behavior. Figure 4 (all transitions enabled) summarizes the behavior model observed at the ARTCC. The switching logic behind the transitions is the object of this section, and is implemented in the form of a cost function, which we describe in this section.

3.1.1 Overall Program Flow

The overall program flow of the simulator is shown in Figure 9. The input to the code is a set of aircraft filed flight plans (Figure 9, middle column), that can either be user generated or taken from ETMS data (as in FACET). As in the true system, these flight plans are not conflict-free and usually do not satisfy metering conditions imposed on the network. Once the program is initialized, aircraft motion simulation follows these flight plans (Figure 9, left column). As time is advanced, conflict as well as metering constraints are dealt with on a sector by sector basis (with sector-wide look ahead, Figure 9, right column), according to the full automaton shown in Figure 4. The flight plans are updated accordingly.

3.1.2 Key Data Structures

Aircraft dynamic equations (1) produce a set of segments; the knowledge of the points connecting the segments and of the aircraft velocity is thus enough to define an aircraft trajectory. This trajectory is thus implemented as a *linked list* of *points* $[x, y, z]$, with a prescribed velocity between the points. The linked list is modified by the simulated controller in the program. The output for each aircraft is the updated linked list. The sectors are implemented as sets of linked lists accessible to a controller. They also contain additional data such as metering conditions (number of aircraft through a given boundary per time unit).

3.1.3 Controller Emulation

ATC behavior is modeled by three levels of priority:

- *Priority 1: No loss of separation.* The prevalent requirement for ATC is to ensure that any aircraft pair is always separated by more than 5 nautical miles.
- *Priority 2: Metering conditions.* The controller needs to ensure that the outflow from his sector is an acceptable inflow for the next sector (or TRACON). Metering conditions can be of various nature: admittance rate or separation at downstream junctions.
- *Priority 3: Best possible throughput.* ATC will try if possible to give direct routes to aircraft in order to minimize their flight times.

These priorities may be modeled using the cost function J:

$$J = \text{cost}_{\text{LOS}} + \text{cost}_{\text{BC breach}} + \text{cost}_{\text{delay}} + \text{cost}_{\text{aircraft actuation}} + \text{cost}_{\text{maneuver}} + \text{cost}_{\text{min dist}} \quad (10)$$

Each term of the cost is a weighted function:

$$J = \sum_{i=1}^{N} \frac{n^i_{\text{LOS}} \cdot w_{\text{LOS}}}{\Delta T^i_{\text{LOS}}} + \sum_{i=2}^{N} (T^i_{\text{breach}})^2 \cdot w_{\text{breach}} + \sum_{i=1}^{N} (TOA^i_{\text{pred}} - TOA^i_{\text{real}}) \cdot w_{\text{delay}}$$
$$+ N_{\text{moved}} \cdot w_{\text{single move}} + \sum_{i=1}^{N} J^i_{\text{maneuver}} + \sum_{i=1}^{N_{\max}} f(d^i_{\min}) \cdot w_{\text{dist}}$$

1. *Loss of separation* (LOS) *cost:* n^i_{LOS} is the number of losses of separation involving aircraft i in the current sector with its current flight plan. ΔT^i_{LOS} is the time until the first loss of separation for aircraft i.
2. *Boundary condition (BC) breach cost:* T^i_{breach} is the time by which an aircraft violates the ΔT time separation constraint from its predecessor (set to zero if the two aircraft are separated by more than ΔT).

3. *Delay cost:* $TOA^i_{\text{pred}} - TOA^i_{\text{real}}$ accounts for the difference between predicted and actual time of arrival (TOA) at the last waypoint of the flight. Positive delays are penalized; earlier arrivals are favored. TOA^i_{pred} and TOA^i_{real} are computed by integration of the flight plans for each aircraft.

4. *Aircraft actuation cost:* N_{moved} accounts for the number of flight plan modifications chosen in the current solution. Large N_{moved} are penalized (the solution chosen by the ATC is often the simplest).

5. *Maneuver cost:* J^i_{maneuver} accounts for the cost of the maneuver selected for aircraft i. Not all maneuvers are of equal preference and therefore have different costs. It is as easy for a controller to prescribe a 10% speed change, a VFS or a shortcut. A holding pattern is the least preferred option, for it requires constant monitoring of the aircraft. This reflects in the weight choice: $J^i_{\text{speed change}} \sim J^i_{\text{shortcut}} \sim J^i_{\text{VFS}} \leq J^i_{\text{holding pattern}}$. The ratio $J^i_{\text{holding pattern}}/J^i_{\text{speed change}}$ is on the order of 10.

6. *Minimal distance cost:* $f(d^i_{\min})$ penalizes aircraft distributions in which aircraft are closely spaced (but do not lose separation) against more sparse distributions. Here, $\text{dist}_{\max} = 7\text{nm}$.
$$\begin{aligned} f(d^i_{\min}) &= \tfrac{1}{d^i_{\min}} \cdot w_{\text{dist}} \quad \text{if } d^i_{\min} < \text{dist}_{\max} \\ f(d^i_{\min}) &= 0 \qquad\qquad\qquad\quad \text{otherwise.} \end{aligned}$$

In order to reflect the three levels of priority of the human ATC stated earlier, the weights shown in the cost function J are: $w_{\text{LOS}} \sim 10^{300} \gg w_{\text{breach}} \sim 10^4 \gg$ other weights ~ 10. Thus a computation which tries to minimize J will first deal with losses of separation, then metering conditions, and finally optimization of the flow. An example of a cost landscape for a given topology of two aircraft is illustrated in Figure 10.

In order to reduce the computational time, we define as N_{choice} the maximum number of aircraft considered for maneuvering by the simulated controller in each time iteration. We restrict N_{choice} to be generally less than the actual number of aircraft per sector. This term is a trade-off between run-time and control-quality and in our simulations it was set in the range of 4 to 8, depending on the targeted goals. Aircraft are selected according to the following rule: aircraft involved in LOS are selected first, then aircraft breaching boundary conditions, and finally remaining aircraft until the selection list has reached N_{choice} aircraft, or until there are no more aircraft to select. In practice, $4 \leq N_{\text{choice}} \leq 8$, where $N_{\text{choice}} = 8$ enables more complicated situations but makes the code run more slowly. The set of all maneuver combinations for the N_{choice} aircraft is called the maneuver set.

At each iteration of the controller emulation loop, an exhaustive recursive search on the maneuver set is run for the chosen aircraft in order to find a set of N_{choice} maneuvers which minimizes J. The computational complexity of finding the optimal J for N_{choice} aircraft subject to n_{maneuver} possible discrete maneuvers is $O((n_{\text{maneuver}})^{N_{\text{choice}}})$. We can reduce this cost to $O((n_{\text{maneuver}} - 2)^{N_{\text{choice}}})$: (i) the cost of the current maneuver has already been computed at the previous step and thus does not need to be recomputed; (ii) two maneuvers

are mutually exclusive (shortcuts), therefore only one needs to be called. Including the cost of checking for conflicts, the total cost of a time iteration becomes:

$$O\left(N_{\max}^2 \cdot (n_{\text{maneuver}} - 2)^{N_{\text{choice}}}\right)$$

where $N_{\max}$ represents the total number of aircraft in the sector. Due to both the discretization of time, and the restriction of the search space to a manageable number of aircraft, our search is not guaranteed to find the global optimum. However, as we show in the next section, the search does provide a reasonable approximation of the controller's behavior. By adjusting the two key control parameters, the number of selected aircraft N_{choice} and time between controller activation ΔT_{act}, a transparent trade-off between run-time and control quality can be made.

3.2 Code validation against ETMS data

The controller logic presented in the previous section is the result of numerous observations we made in the Oakland ARTCC, monitoring the work of air traffic controllers. The fact that one can classify ATC action into a set of preferred directives was experimentally validated for a different airspace (see Hansman and Histon[20]). However, even if the automaton of Figure 4 and the cost function of the previous section implemented in the simulator are consistent with our observations, there is no a priori guarantee that these would produce the same effects on the system as a human controller. For this reason, we assess how well our cost function describes the decision making of a human controller by validating our code against recorded aircraft trajectories. We use ETMS data provided by NASA Ames.[2]

We first explain the data extraction process which enables us to convert ETMS data to a readable format for our simulator. We then present the three types of validations realized.

3.2.1 Data extraction

We can extract two types of information from the ETMS data: the actual flown aircraft trajectories, and the filed flight plans for each aircraft (eventually updated if modifications are made during the flight). The position is given in latitude / longitude and in terms of navaids, fixes, and jetways which we represent in cartesian coordinates, an approximation valid for the portion of airspace of interest to us. The filed flight plan is given in terms of navaids, fixes, jetways, which we also convert into cartesian coordinates using a public database (`http://www.airnav.com`). Future versions of our simulator might use recently developed ETMS analyses tools such as [13] to perform these tasks automatically. This study

[2]Data is collected from the entire population of flights with filed flight plans in the NAS. ETMS data is sent from the Volpe National Transportation System Center (VNTSC) to registered participants via the Aircraft Situation Display to Industry (ASDI) electronic file server. A file containing all recorded data is generated. It displays for each aircraft the current flight data (time, position, speed, heading), as well as the filed flight plan (in terms of navaids, jetways, fixes, etc.). The update rate of the measurements is of the order of one minute.

is limited to sector control, and we did not implement the Traffic Management Unit (TMU) action. TMU operates at the Center level and makes strategic flow scheduling decisions, which go beyond the range of a single sector controller. We therefore need to validate our simulator at a scale at which TMU actuation is already incorporated in the flight plans (typically one or two sectors). Since our interest focuses on sectors 32, 33, 34, 15, and 13, the actual flight plans are truncated, and we keep only points in that sector. The filed flight plans are truncated similarly. Their estimated time of entrance in the sector is set to the actual time of entrance of the closest actual recorded point. The entrance point in the computational domain is taken to be the closest conflict-free point to the intersection of the flight plan and the boundary of the corresponding sector. The altitude assigned to the filed flight plans is set to the average altitude of the actual trajectory in that sector.

3.2.2 Validation

Comparison of flight times. We select the first 100 aircraft of the ETMS data set which are flying above 33000ft for more than 6 minutes. Their recorded trajectories are extracted as sequences of waypoints which are used as initialized flight plans for our simulations. We then compare the flight times in the simulation to the actual ones. The experiment is run for the following set of speeds: $M \in \{0.6, 0.7, 0.8\}$ (M is the Mach Number), which matches our observations in the data for this altitude. In the run, the controller is activated every $\Delta T_{\text{act}} = 10$sec. The results are shown in Figure 11. Two main conclusions can be made: *(i)* the simulator is able to recreate the flow without major modification, and eventually resolves apparent conflicts in the data – these conflicts can be due to inaccuracy of the measurements or transmission (two aircraft separated by less than 5 nm at the same altitude), or due to problems of interpolation when speed changes in time; *(ii)* the time comparison (Figure 11) shows relatively good matching. The flight times provided by the simulator are usually shorter than in reality because by default the simulator will always try to maximize the throughput in the sector. The mean deviation is 120 sec. for flight plans with an average duration of 1300 sec. (9.2% error).

Validation of conflict resolution. We select aircraft flying through sector 33 in a time frame of 10 hours (a total set of 314 aircraft) and simulate these flights. The filed flight plans are not conflict-free. We would like to show that the simulator is able to provide a conflict-free environment. For this run, the activation time of the controller is $\Delta T_{\text{act}} = 20$sec.[3] The set of speeds allowed is $M \in \{0.55,\ 0.75,\ 0.89\}$ (since we are considering the full range of altitude, we need to consider the full range of speed). The simulator is able to provide a conflict free environment. During the simulation, it has to actuate 50 different aircraft. The number of resolved conflicts can be assumed to be on the same order, since a single intervention will usually resolve not more than one conflict.

Validation of maneuver assignments. The core of the simulator is the model of the human

[3] $\Delta T_{\text{act}} = 10$sec. or $\Delta T_{\text{act}} = 20$sec. is on the order of the maximal actuation rate of a controller. We choose $\Delta T_{\text{act}} = 20$sec. in this particular case because of the duration of the computation (10 hours of real time are simulated).

controller by a decision procedure based on the cost function described previously. The validation so far has shown the correlation of flow patterns generated by our code and these observed in reality. We would also like to assess the validity of our decision procedure. For this, we identify conflict resolution maneuvers, which are typically obtained by identifying deviations from the filed flight plan. For these maneuvers, we generate the following input data: all aircraft are assigned their actual flight plan (recorded trajectories), and the aircraft for which the maneuver is identified is assigned its filed flight plan (a set of way points). We thus put the simulator in the same situation as the human controller, in which it has to make the decision that was actually taken. For sector 33, we were able to identify 20 distinct maneuvers out of 300 examined flight plans. The simulator reproduced 16 of them.[4] An example is shown in Figure 12.

3.3 Validation of the analytical predictions

In this section, we assess how accurate our analytical predictions are with respect to the real system, by comparing them with simulations. We present an example of two backpropagating shocks, solved with an extension of the method explained in Section 2. Two platoons, each at 10 miles-in-trail[5], are respectively subjected to 15 miles-in-trail and 20 miles-in-trail outflow boundary conditions (the boundary conditions for platoon 2 start after all aircraft of platoon 1 have reached the TRACON at time $t = 4300$). The speeds are $M \in \{0.59,\ 0.75,\ 0.89\}$. The resulting aircraft flows for this analytic solution are shown in Figure 13. The results and interpretations are shown in Figures 15 and 14. Two shocks appear successively. From Figure 15, we can see that within the second platoon, the first twelve aircraft need to be actuated within the Oakland Center, whereas the last eight need to be actuated upstream (Salt Lake Center). Since in general, no knowledge of the required boundary conditions is propagated upstream, we can predict that in the real system, the last eight aircraft would not be actuated until they enter the Oakland Center and that no solution to this metering problem would be found without putting the aircraft on hold.

We verify this by simulating this flow. In Figure 14, we see that for the last eight aircraft, the first activation time in the simulator is higher than the predicted (upper plot): the simulated controller is not able to actuate the aircraft on time, because they are not in its airspace. We illustrate in the middle plot that these aircraft are breaching the boundary conditions (by about one minute each), which can also be seen in the lower plot: their delays (the inverse of the cumulated breaches) become negative, i.e. they arrive in advance. This is an illustration of distributed and decentralized control: the actuation occurs in different sectors, and the only communication between the sectors is through the metering

[4]Small-scale maneuvers are less likely to be executed correctly by the simulator because the probability of selecting the respective aircraft at exactly the 'right' time is small, which explains the small discrepancy between the results. Also, even if the maneuver is executed correctly by the simulator, the resulting flight plan will look different from the ETMS data, since the simulator is restricted to a single angle of deviation ($\theta = 22.5°$).

[5]The terminology "n miles-in-trail" is a standard term used by ATC. It means that aircraft follow each other separated by n miles.

conditions. Obviously, the lack of centralized actuation (here communication and strategic TMU planning) disables efficient flow scheduling.

3.4 Decorrelation of flow metering and conflict resolution

A working assumption which is often made in flow scheduling is the decorrelation of metering conditions and cross traffic when traffic density is relatively low, that is, that conflict avoidance actuation does not impact metering actuation. We demonstrate this property of the flow by simulating two streams of intersecting (and conflicting) aircraft (see Figure 16). Two platoons of aircraft intersect at a navaid (Clovis, in Sector 15), and we consider pairwise conflicts. One platoon is subject to outflow metering conditions (at the boundary of sector 15 and 34), and we would like to quantify the impact of the other platoon on the travel time of the first platoon through the sector.

We investigate 16 different configurations by selecting the same inflow conditions for the two platoons: $\Delta L_{\text{in}} \in \{15, 20, 25, 30\}$ miles-in-trail, and the following outflow conditions for the platoon heading towards sector 34, $\Delta L_{\text{out}} \in \{15, 20, 25, 30\}$ miles-in-trail (see Figure 16). We compare the travel time for each platoon without the presence of the other, with the travel time when the two platoons intersect. For each configuration, we perform 10 runs where the initial position of the aircraft within each platoon is perturbed by a uniform noise of amplitude 2 nm. This scenario represents ths situation in which each aircraft from one platoon conflicts pairwise with an aircraft from the second platoon. This results in 160 runs.[6]

For each $(\Delta L_{\text{in}}, \Delta L_{\text{out}})$, we compute the three following quantities

$$\begin{array}{lrcl} 1. & \delta T^{\text{no cross flow}}_{\text{travel time}} & = & \frac{1}{N}\sum_{i=1}^{N}\left(T^{\text{no cross flow, B.C.}}_{\text{aircraft } i} - T^{\text{no cross flow, no B.C.}}_{\text{aircraft } i}\right) \\ 2. & \delta T^{\text{cross flow}}_{\text{travel time}} & = & \frac{1}{N}\sum_{i=1}^{N}\left(T^{\text{cross flow, B.C.}}_{\text{aircraft } i} - T^{\text{no cross flow, no B.C.}}_{\text{aircraft } i}\right) \\ 3. & \Delta T_{\text{due to cross flow}} & = & \delta T^{\text{cross flow}}_{\text{travel time}} - \delta T^{\text{no cross flow}}_{\text{travel time}} \end{array} \tag{11}$$

Here, N is the total number of aircraft ($N = 20$, we have two platoons of 10 aircraft). In each formula above, $T^{\text{no cross flow, B.C.}}_{\text{aircraft } i}$ represents the travel time of aircraft i in absence of the other platoon, and with outflow boundary conditions as shown in Figure 16 (similar definitions apply for the other variables in formula (11) and are obvious from notation). The results for the mean difference in travel time are shown in Figure 17 (averaged over 10 runs for each case). We show the numerical results obtained respectively in the $\Delta L_{\text{in}} = 15$ miles-in-trail inflow case:

ΔL_{out}	$\delta T^{\text{no cross flow}}_{\text{travel time}}$	$\delta T^{\text{cross flow}}_{\text{travel time}}$	$\Delta T_{\text{due to cross flow}}$
15nm	0.7sec.	140.5sec.	139.8sec.
20nm	64.8sec.	166.9sec.	102.1sec.
25nm	90.7sec.	157.0sec.	66.3sec.
30nm	104.6sec.	164.8sec.	60.2sec.

(12)

[6]The settings for these runs are: $M \in \{0.8,\ 0.85,\ 0.89\}$, the vector-for-spacing maneuver was limited to a 5 nm deviation from the original flight plan. These settings were chosen to guarantee short flight times. The interval between controller activation was set to $\Delta T_{\text{act}} = 20sec.$

and for the $\Delta L_{\text{in}} = 20$ miles-in-trail inflow case:

ΔL_{out}	$\delta T^{\text{no cross flow}}_{\text{travel time}}$	$\delta T^{\text{cross flow}}_{\text{travel time}}$	$\Delta T_{\substack{\text{due to}\\ \text{cross flow}}}$
15nm	0sec.	28.1sec.	28.1sec.
20nm	0.7sec.	35.9sec.	35.2sec.
25nm	74.2sec.	101.9sec.	27.6sec.
30nm	99.5sec.	123.9sec.	24.3sec.

(13)

$\Delta T_{\substack{\text{due to}\\ \text{cross flow}}}$ is represented in Figure 17 for the complete set of $(\Delta L_{\text{in}}, \Delta L_{\text{out}})$ investigated here. Even if the peak of $\Delta T_{\substack{\text{due to}\\ \text{cross flow}}}$ happens for $(\Delta L_{\text{in}}, \Delta L_{\text{out}}) = (15, 15)$, the maximum $\delta T^{\text{cross flow}}_{\text{travel time}}$ happens as expected for $(\Delta L_{\text{in}}, \Delta L_{\text{out}}) = (15, 30)$, which is the maximal inflow, minimal outflow, as can be seen in (12). The comparison between (12) and (13) and Figure 17 clearly shows the predominance of conflict resolution over boundary conditions for high density of traffic (see last column in (12) and (13)).

We see that the difference in delay between separate and simultaneous flow is significantly larger if the aircraft are spaced at 15 nm when compared to platoons with larger spacing (see Figure 17). While the mean difference per plane is always larger than 60 seconds for the 15 nm-platoons, it is always smaller than 60 seconds for platoons with larger spacing. With an average flight time of 660 seconds over all scenarios, 60 seconds corresponds to an average delay of 9% in flight time. The worst case difference (15 nm inflow, 30 nm outflow) is more than 21% of the overall flight time. These numbers are significant, especially when considering the possibility of multiple intersecting platoons.

Note that in the left plot of Figure 17, we would intuitively expect the largest difference in delay (cross flow vs. no cross flow) to happen for $(\Delta L_{\text{in}}, \Delta L_{\text{out}}) = (15, 30)$, which is the hardest situation to achieve (maximum inflow, most restricted outflow). In fact, this maximum occurs at $(\Delta L_{\text{in}}, \Delta L_{\text{out}}) = (15, 15)$, which does not seem a priori to be problematic (conservation of flow). This can be explained by looking at the middle and right plot of Figure 17. In the absence of cross flow, the delay accumulated due to the boundary conditions is maximal for $(\Delta L_{\text{in}}, \Delta L_{\text{out}}) = (15, 30)$, as expected (right plot of Figure 17).When subtracted from the difference in travel time with cross flow (middle plot in Figure 17), a max appears at $(\Delta L_{\text{in}}, \Delta L_{\text{out}}) = (15, 15)$, because $\delta T^{\text{no cross flow}}_{\text{travel time}}|_{(\Delta L_{\text{in}}, \Delta L_{\text{out}})=(15,15)} = 0$. Thus, these results are consistent. The flows of Figure 16 are worst cases, and it is rare to observe such a density in both directions. However, sectors such as 33 are often subject to less structured, but just as dense cross flows, and our predictions are relevant for this type of airspace.

Currently, TMU does not take the influence of conflict resolution into account when making decisions, since the impact of ATC on the sector level is considered to be negligible for overall flight times. In the current system this is not always true, thus leading to inaccuracies in the predictions of sector occupancy, i.e. how many aircraft will be in a sector in $t + 15$ minutes. We have thus shown in a particular case that the influence of conflict resolution actuation increases with higher traffic density, which therefore necessitates information feedback from the sector to the Center.

4 Conclusion and Current Work

We have derived a control theoretic model of sector-based traffic flow using hybrid automata theory. We used a subset of this model to generate Lagrangian analytic predictions of the traffic flow (dynamic sector capacity, extend of traffic jams), which were linked to Eulerian models of the NAS. We used these results to predict the conditions under which airspace saturation cannot be treated at the level of a single sector, but requires centralized actuation (communication and strategic TMU planning). These predictions were validated against an abstraction of the real system (a simulator using the full model, which we validated against real data). Finally, we generated flow conditions under which we can decorrelate metering from conflict resolution.

This last result is useful for our current work, because it enables us to ignore conflict resolution for scheduling. Our current work [23] aims at generating fast algorithms to design metering protocols for this converging airspace. The emphasis of this work is in design, rather than modeling, and it uses all modes of the automaton shown in Figure 4. Therefore, no analytical solution is available to solve this problem. In fact, the linear program derived in section 2 generalizes to a Mixed Integer Linear Program (MILP), which we believe to be NP-complete in the general case. However, in specific cases relevant to ATC automation, we have been able to derive polynomial time algorithms to solve this program, which leads us to believe that this method could be applied to real time scheduling and airspace decongestion.

Acknowledgments

We are grateful Gano Chatterji for his initial suggestions concerning the construction of the shock wave, to Banavar Sridhar for fruitful discussions about modeling this problem, Shon Grabbe for his help on FACET and ETMS data, and Karl Bilimoria for his recommendations concerning the converging flow.

References

[1] M. S. Nolan, *Fundamentals of Air Traffic Control, 3rd Edition*. Pacific Grove, CA: Brooks/Cole Publishing, 1999.

[2] T. S. Perry, "In search of the future of air traffic control," *IEEE Spectrum*, vol. 34, no. 8, pp. 18–35, 1997.

[3] J. Jackson and S. Green, "Control applications and challenges in air traffic management," in *Proceedings of the 1998 American Control Conference*, (Philadelphia, PA), 1998.

[4] S. Kahne and I. Frolow, "Air traffic management: Evolution with technology," *IEEE Control Systems Magazine*, vol. 16, no. 4, pp. 12–21, 1996.

[5] R. Y. Gazit, *Aircraft Surveillance and Collision Avoidance using GPS*. PhD thesis, Department of Aeronautics and Astronautics, Stanford University, 1996.

[6] N. Pujet and E. Feron, "Flight plan optimization in flexible air traffic environments," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (San Diego, CA), August 1996.

[7] J. Plaettner-Hochwarth, Y. Zhao, and J. Robinson, "A modularized approach for comprehensive air traffic system simulation," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Denver, CO), August 2000.

[8] H. Erzberger, T. J. Davis, and S. Green, "Design of Center-TRACON Automation System," in *Proceedings of the AGARD Guidance and Control Symposium on Machine Intelligence in Air Traffic Management*, (Berlin, Germany), pp. 11.1–11.12, 1993.

[9] A. Odoni, J. Bowman, D. Delahaye, J. Deyst, E. Feron, J. Hansman, K. Khan, J. Kuchar, N. Pujet, and R. Simpson, "Existing and required modeling capabilities for evaluating ATM systems and concepts," Tech. Rep. NAG2-997, MIT, 1996.

[10] J. Kuchar and L. Yang, "A review of conflict detection and resolution modeling methods," *IEEE Transactions on Intelligent Transportation Systems*, vol. 1, no. 4, 2000.

[11] K. Bilimoria, B. Sridhar, G. Chatterji, K. Seth, and S.Graabe, "FACET: Future ATM concepts evaluation tool," in *Proceedings of the 3rd USA/Europe Air Traffic Management R&D Seminar*, (Naples, Italy), June 2001.

[12] K. Bilimoria and H. Lee, "Properties of air traffic conflicts for free and structured routing," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Montreal, PQ), August 2001.

[13] J. C. Lin and D. K. Gifford, "VisualFlight: The air traffic control data analysis system," tech. rep., MIT, May 2002.

[14] D. Dugail, E. Feron, and K. Bilimoria, "Conflict-free conformance to En-Route flow-rate constraints," in *Proceedings of the AIAA Conference on Guidance, Navigation and Control*, (Monterey, CA), August 2002.

[15] P. Mishra and G. Pappas, "Flying hot potatoes," in *Proceedings of the 2002 American Control Conference*, (Anchorage, AK), pp. 754–759, May 2002.

[16] M. Innocenti and L. Pollini, "Spatial trajectory generation for conflict avoidance in air traffic management," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Denver, CO), August 2000.

[17] Z. Mao, E. Feron, and K. Bilimoria, "Stability of intersecting aircraft flows under decentralized conflict avoidance rules," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Denver,CO), August 2000.

[18] J. Krozel, M. Peters, and K. Bilimoria, "A decentralized control strategy for distributed air/ground traffic separation," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Denver,CO), August 2000.

[19] P. K. Menon, G. Sweriduk, and K. Bilimoria, "A new approach for modeling, analysis and control of air traffic flow," in *Proceedings of the AIAA Conference on Guidance, Navigation and Control*, (Monterey, CA), August 2002.

[20] J. Histon and R. J. Hansman, "The impact of structure on cognitive complexity in air traffic control," Tech. Rep. ICAT-2002-4, MIT, June 2002.

[21] B. Sridhar, K. S. Sheth, and S. Grabbe, "A new approach for modeling, analysis and control of air traffic flow," in *Proceedings of the* 2^{nd} *USA/Europe Air Traffic Management R&D Seminar*, (Orlando, FL), December 1998.

[22] D. Bertsimas and S. Stock Patterson, "The air traffic flow management problem with enroute capacities," *Operations Research*, vol. 46, pp. 406–422, 1998.

[23] A. M. Bayen and C. J. Tomlin, "Real-time discrete control law synthesis for hybrid systems using MILP: applications to congested airspaces." Submitted to the *2003 American Control Conference*.

[24] Jeppesen, "High altitude enroute charts." http://www.jeppesen.com, February 2000.

[25] D. Dugail, Z.-H. Mao, , and E. Feron, "Stability of intersecting aircraft flows under centralized and decentralized conflict avoidance rules," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Montreal), August 2001.

[26] L. Pallottino, A. Bicchi, and E. Feron, "Mixed integer programming for aircraft conflict resolution," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Montreal), August 2001.

[27] K. Bilimoria, "A geometric optimization approach to aircraft conflict resolution," in *Proceedings of the AIAA Guidance, Navigation, and Control Conference*, (Denver,CO), August 2000.

[28] A. Nilim, L. El-Ghaoui, V. Duong, and M. Hansen, "Trajectory-based air traffic management (TB-ATM) under weather uncertainty," in *Proceedings of the 4rd USA/Europe Air Traffic Management R&D Seminar*, (Santa Fe, NM), December 2001.

[29] A. M. Bayen, P. Grieder, H. Sipma, G. Meyer, and C. J. Tomlin, "Delay predictive models of the National Airspace System using hybrid control theory," in *Proceedings of the 2002 American Control Conference*, (Anchorage, AK), May 2002.

[30] A. M. Bayen, P. Grieder, and C. J. Tomlin, "A control theoretic predictive model for sector-based air traffic flow," in *Proceedings of the AIAA Conference on Guidance, Navigation and Control*, (Monterey, CA), August 2002.

[31] G. F. Newell, "A simplified theory of kinematic waves in highway traffic, part I: General theory," *Transportation Research*, vol. 27, no. 4, pp. 281–287, 1993.

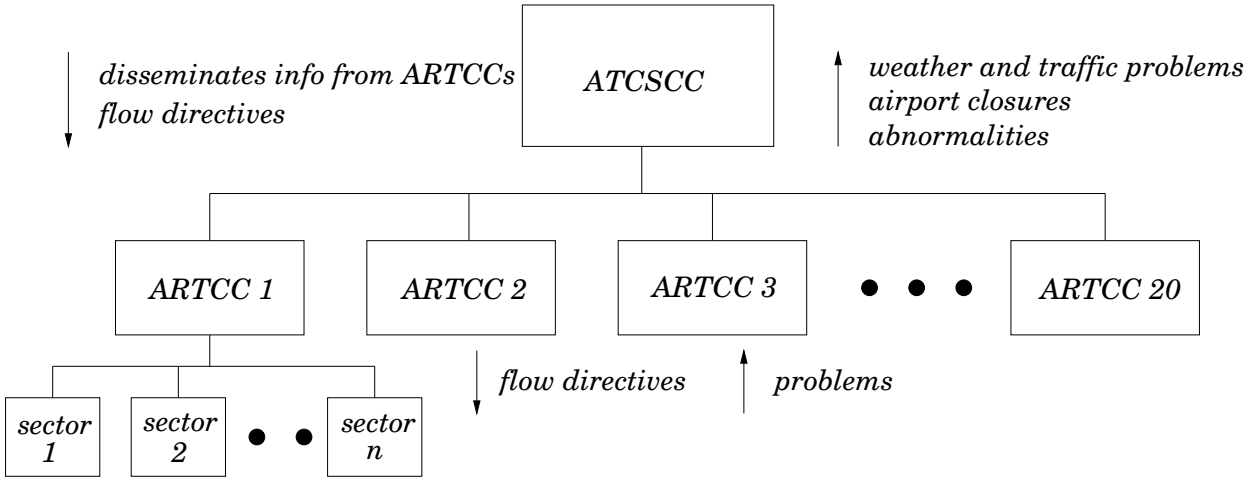

Figure 1: Control hierarchy in the current structure of NAS.

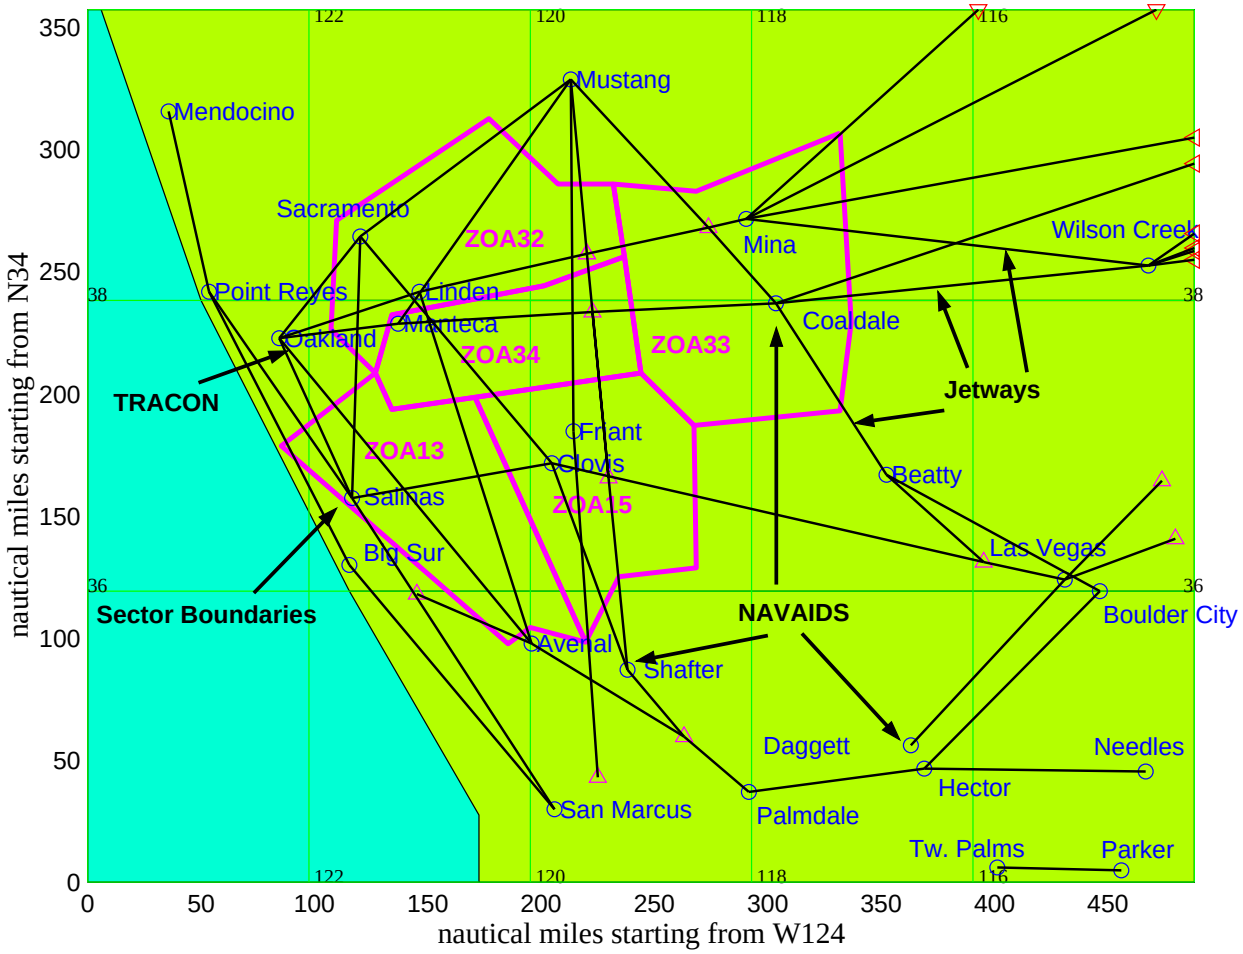

Figure 2: ATC sectors modeled for this study: 32, 33, 34, 13 and 15 within the Oakland ARTCC (labeled above as ZOA32, etc.). Most crucial jetways and waypoints for a San Francisco approach are shown here. The data modeled comes from FACET [11] as well as JEPPESEN high altitude en route charts [24].

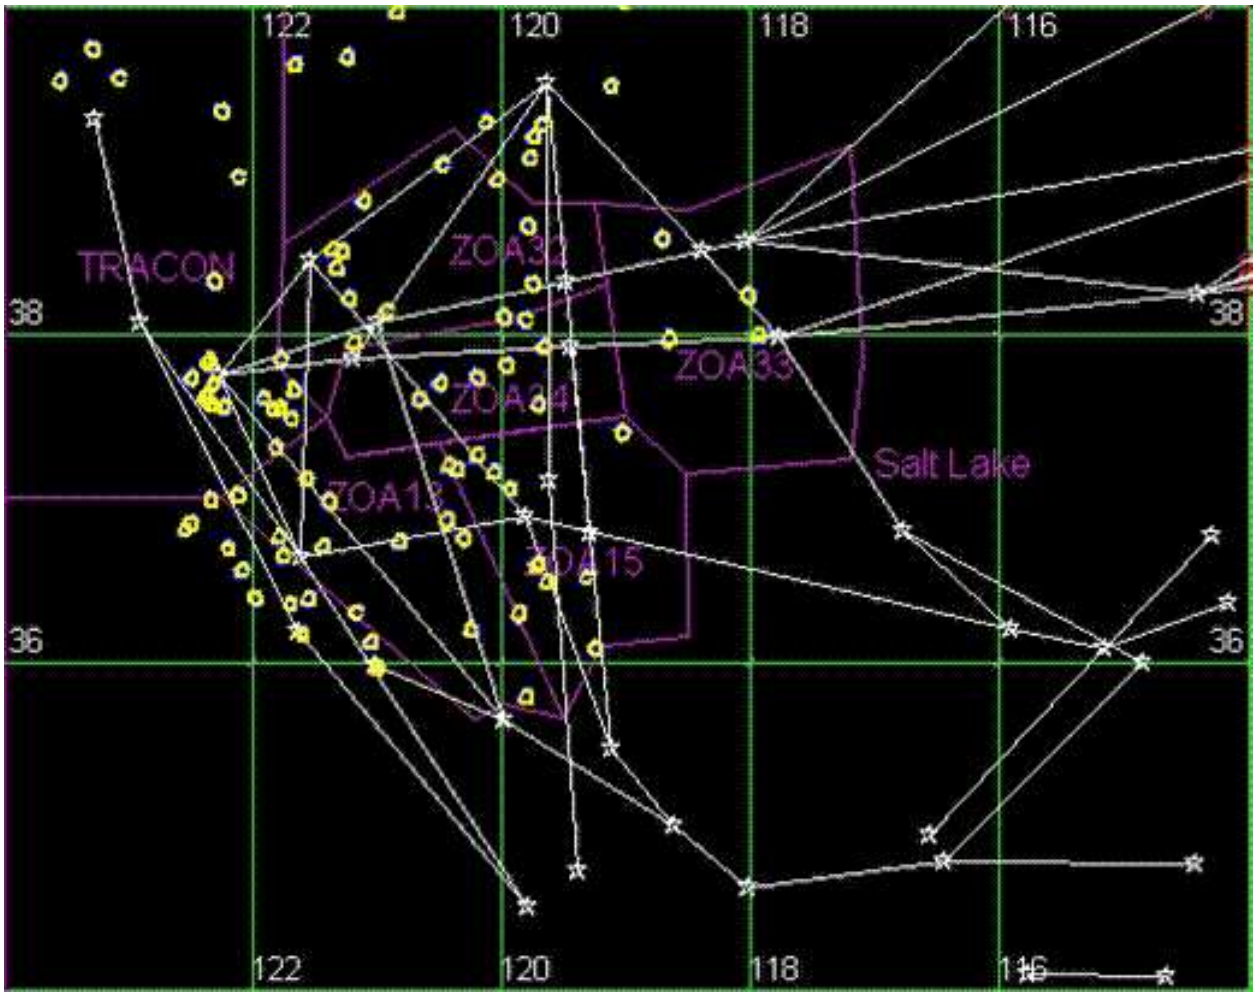

Figure 3: Visual display of the simulator. Traffic in the Oakland ARTCC (aircraft not in this Center have been filtered out). This plot has been generated using ETMS data.

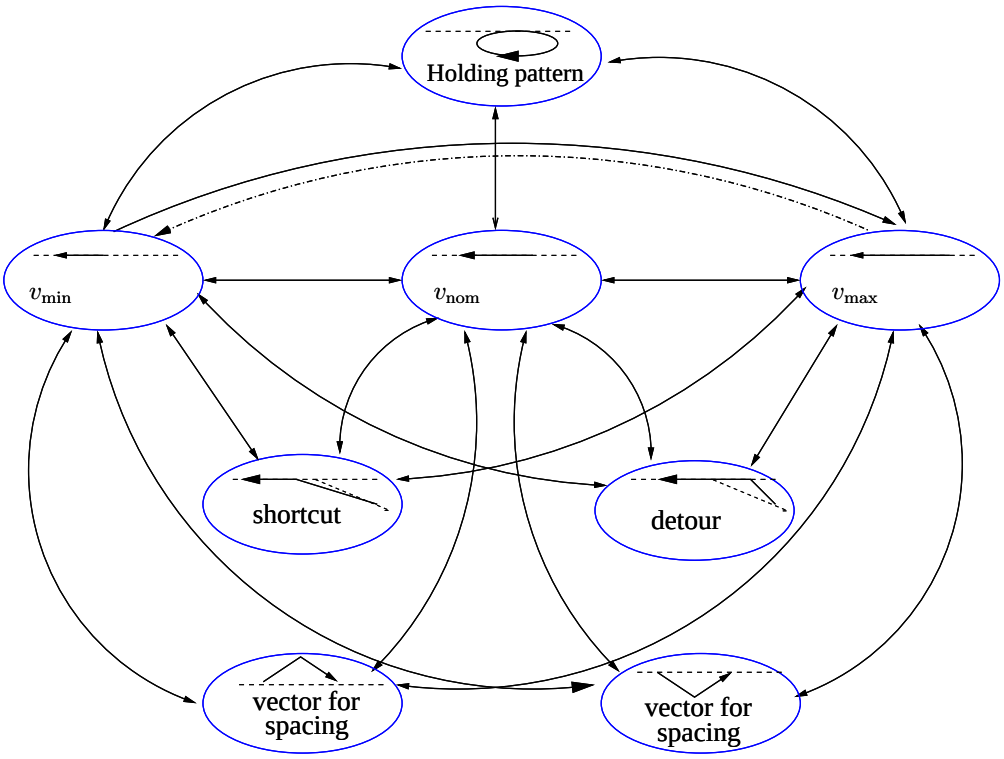

Figure 4: Hybrid automaton representing the action of one controller on a single aircraft. Each of the eight modes represents one possible state of the aircraft. The arrows joining these states are the mode switches, initiated by the controller. The dash-dotted transitions are used for the analytical solution. The complete set of arrows is used for the simulation.

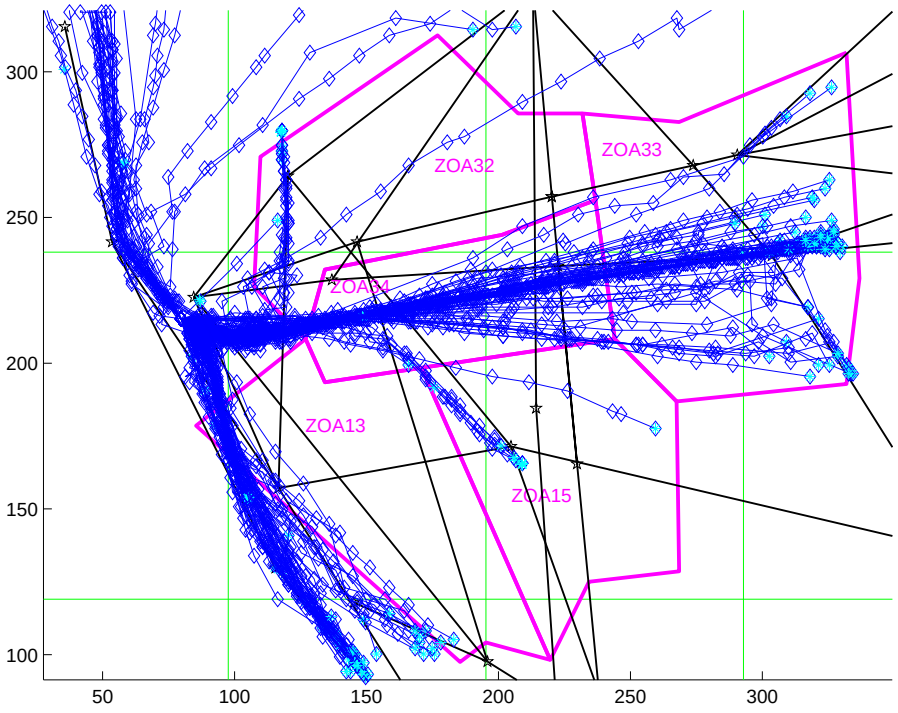

Figure 5: Overlay of trajectories merging into San Francisco (11 hours of traffic). The data modeled comes from ETMS and FACET [11]. Our Lagrangian approach models each of the trajectories included in this plot.

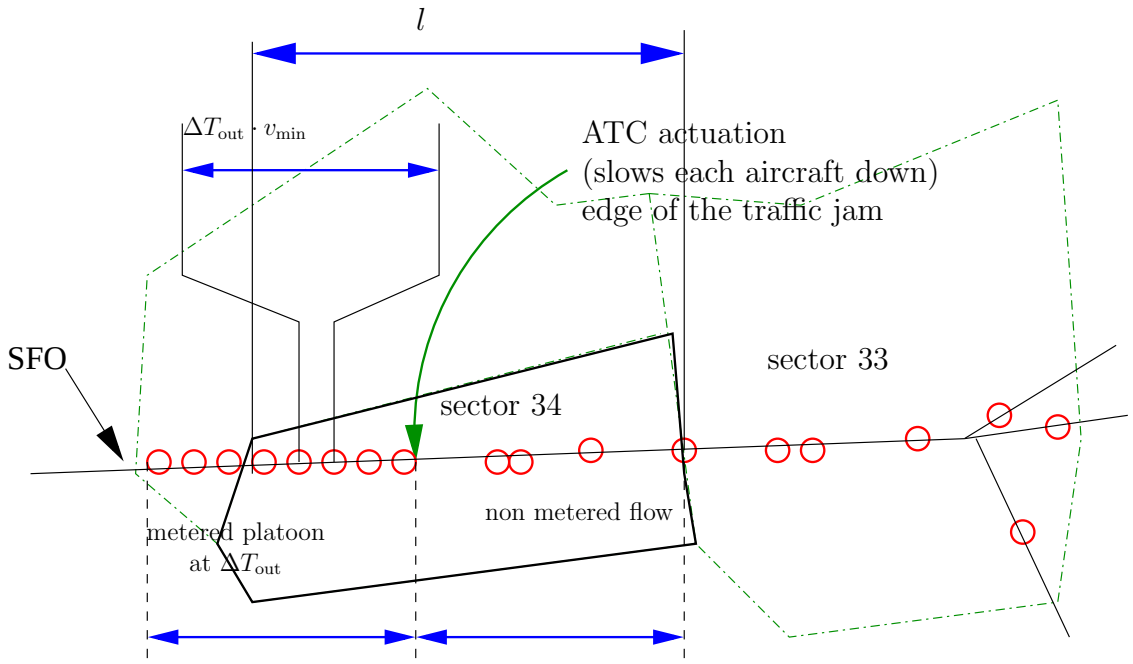

Figure 6: ATC control action on the merging flow. The traffic jam extends from SFO to the edge of the platoon with $\Delta T_{\text{out}} \cdot v_{\min}$ miles-in-trail (one aircraft every ΔT_{out} time units). Once the actuation point (edge of the traffic jam) has moved upstream (into sector 33), sector 34 is called saturated.

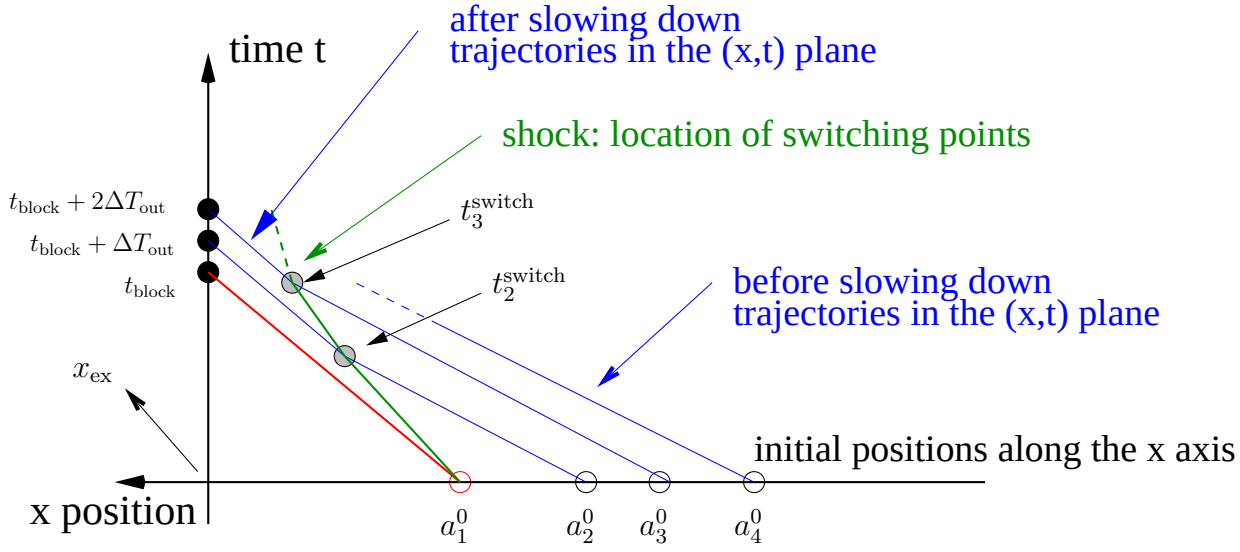

Figure 7: Shock construction. The aircraft trajectories are represented in the (x, t) plane. They originate at $t = 0$ from the horizontal axis (white circle on each trajectory). After some amount of time, the aircraft may be switched to speed $v_{\min}$ at location $(x_i^{\text{switch}}, t_i^{\text{switch}})$ (shaded circle on each trajectory). Ultimately they reach x_{ex}, the entrance of TRACON (black circle).

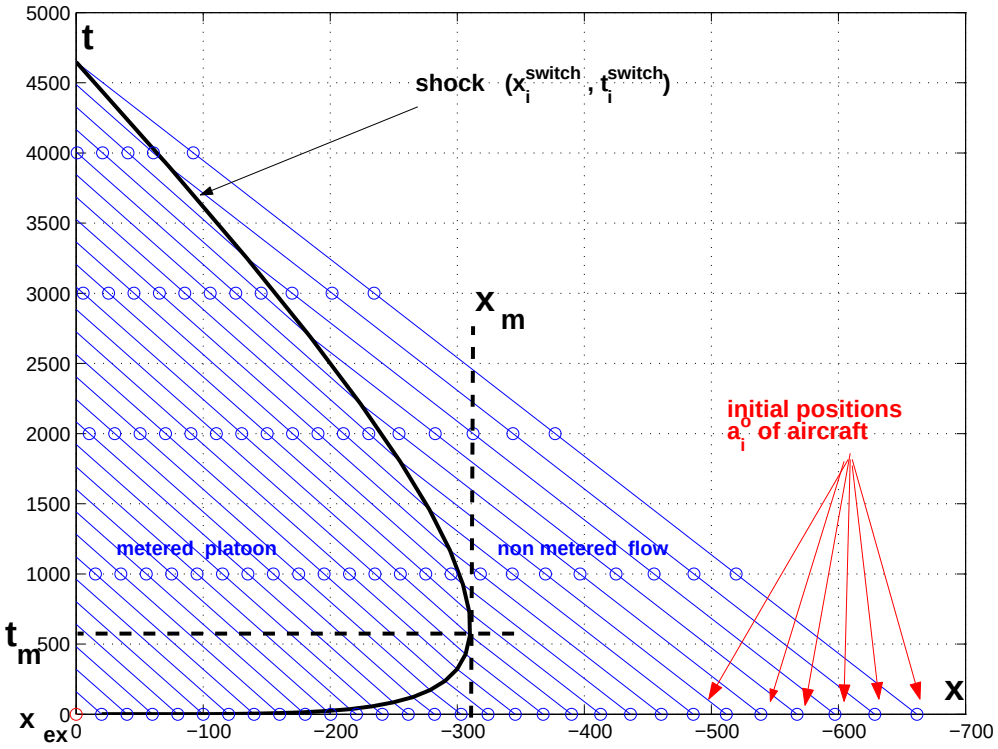

Figure 8: Switching curve (shock) for a vanishing traffic jam. x denotes the distance to the metering point (SFO). The lines are the trajectories of the aircraft in the (x, t) space. The positions of aircraft are represented every 1000 sec. as dots. Once they have passed through the shock, they are separated by $v_{\min}\Delta T_{\text{out}}$. The point (x_m, t_m) is the furthest reachable point by this traffic jam. Note that the slope of the lines changes through the shock. The slope difference can hardly be seen visually, because the speed change is small.

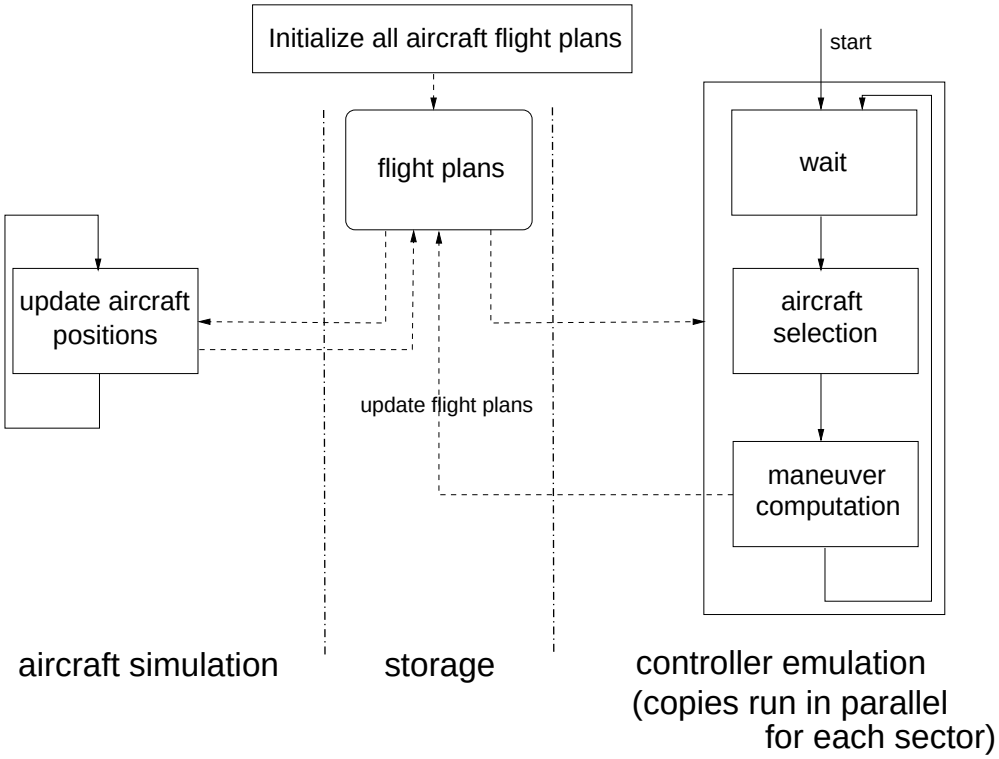

Figure 9: Program flow of the simulator.

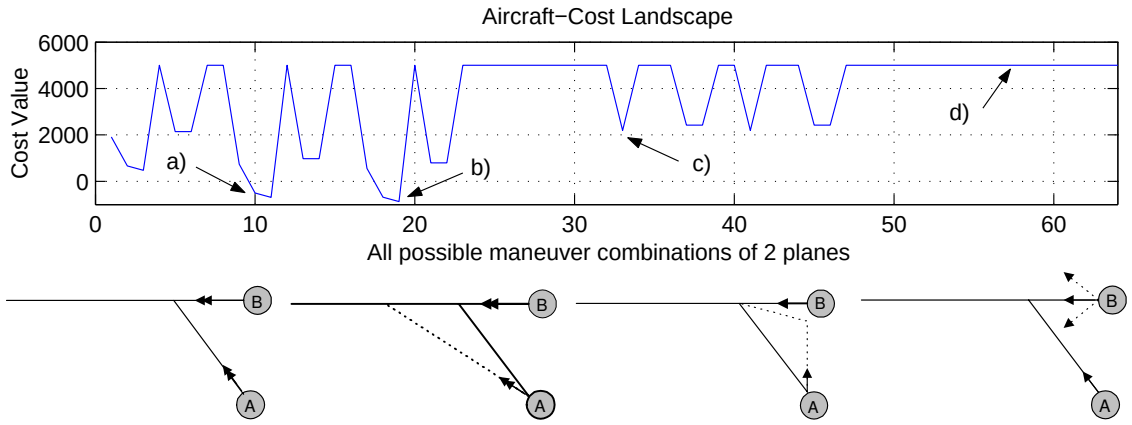

Figure 10: Top: cost values for all possible maneuver combinations in a two-aircraft intersection scenario, where the eight maneuvers of Figure 4 are enabled (thus generating $8^2 = 64$ possible values of J). Four out of 64 examples are extracted and illustrated (four lower pictures). (a) Both aircraft maintain same speed; (b) Aircraft A takes a shortcut maintaining aircraft B at max speed; (c) A makes a VFS at low speed; (d) A does nothing, B is not able to prevent the loss of separation. In this case, the simulated controller would choose solution (b) since the lowest cost is associated with that maneuver. The cost J has been truncated at $5 \cdot 10^3$ for readability. Thus a LOS cannot be visually differentiated from a breach in this plot, though it can in our data. Thus, too large breaches as well as LOS do not appear on this plot.

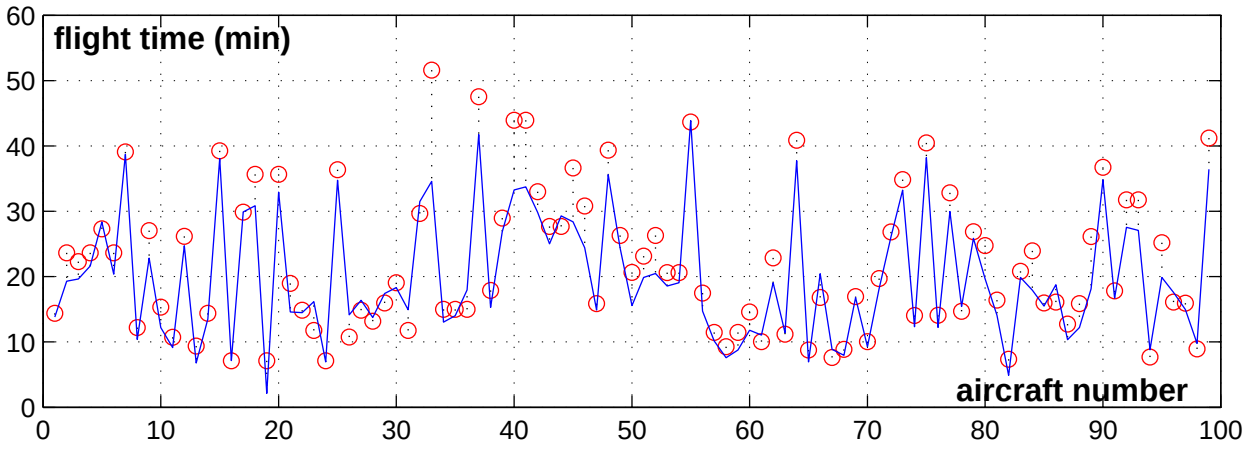

Figure 11: Flight time comparisons for the first 100 aircraft going through sector 33 in the ETMS data set we used. The dots are the flight times for the ETMS recorded points. The solid curve is the result of our simulations.

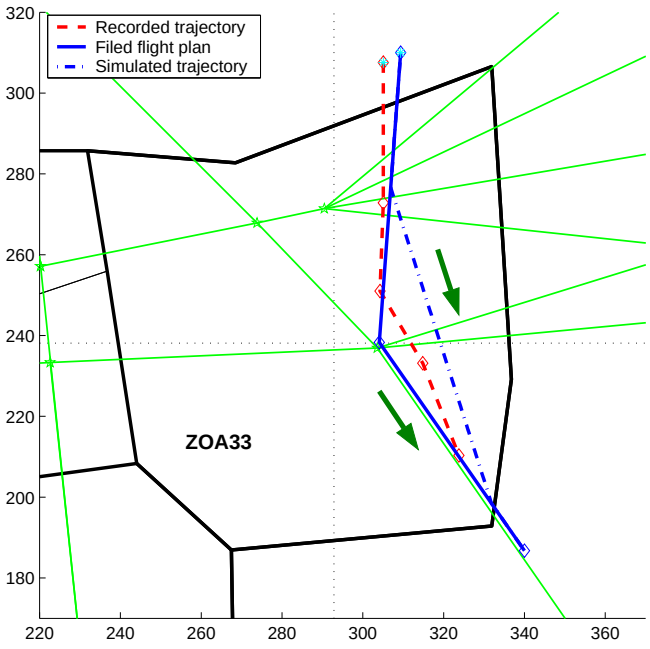

Figure 12: An example of maneuver caused by conflict resolution, reproduced by the simulator. The recorded data (dashed) exhibits an actual shortcut from the filed flight plan (solid). The simulated trajectory (dashed-dotted) is a shortcut of the same type.

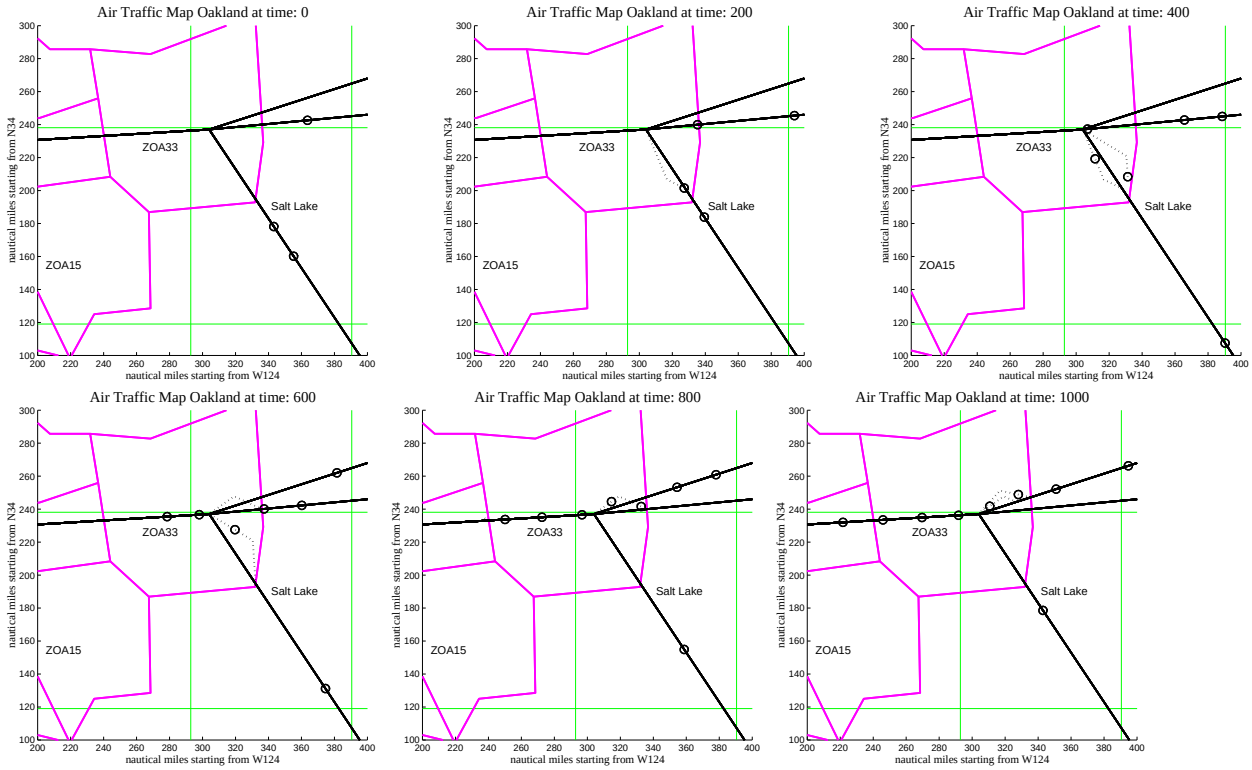

Figure 13: Sector 33: traffic flow for the merging traffic simulation of Figures 15 and 14. The radius around the aircraft is 2.5 nm. The solid lines represent the aircrafts' flight plan. The dotted lines correspond to maneuvers assigned by the simulator. The simulator makes extensive use of the vector-for-spacing to meter the aircraft.

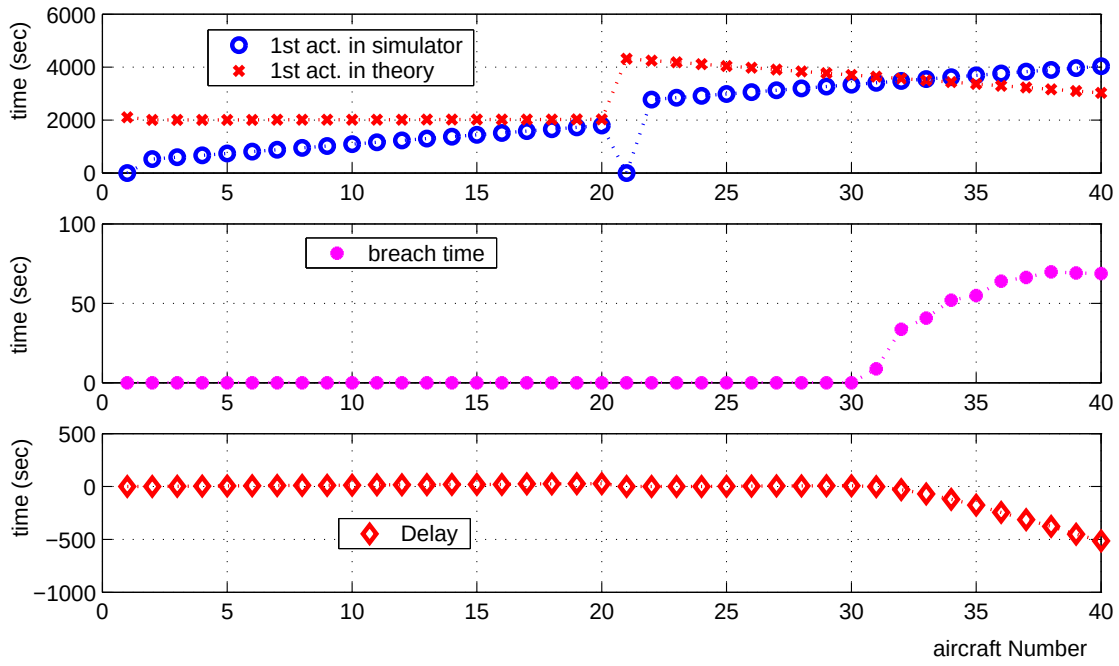

Figure 14: First actuation times of the aircraft (upper plot), simulated and predicted; breaches in metering conditions (middle plot), simulated; delays (lower plot), simulated, for the case of the two platoons of Figure 15.

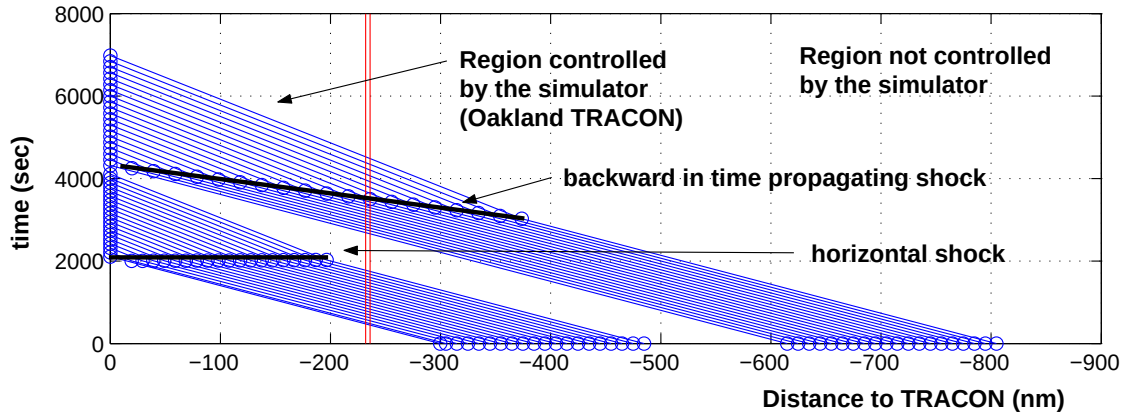

Figure 15: Shocks generated by two successive platoons. The first shock is steady in time (it only propagates backward in space). It corresponds to a piling up process on a highway where all vehicles slow down at the same time. The second shock propagates backward in space and time (which is much harder to handle in practice, because actuation must be performed upstream first).

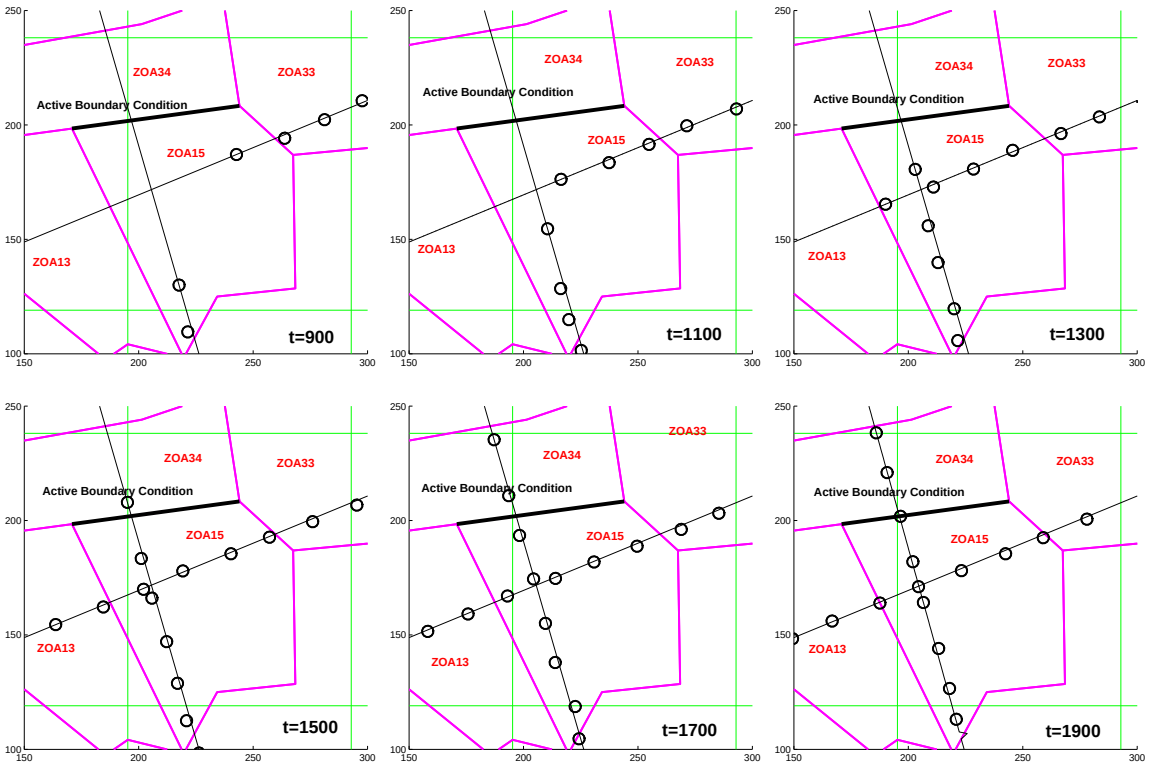

Figure 16: Two intersecting platoons at the Clovis navaid for the configuration in which the aircraft have an inflow of 20 miles-in-trail. Aircraft are pairwise conflicting (aircraft i from platoon 1 conflicts with aircraft i of platoon 2), assuming the flight plan is to follow the straight line at the initial speed. The simulator is thus forced to adjust each of the pairs. Deviations (vector-for-spacing) are barely visible on the plot because of their small amplitude (5nm).

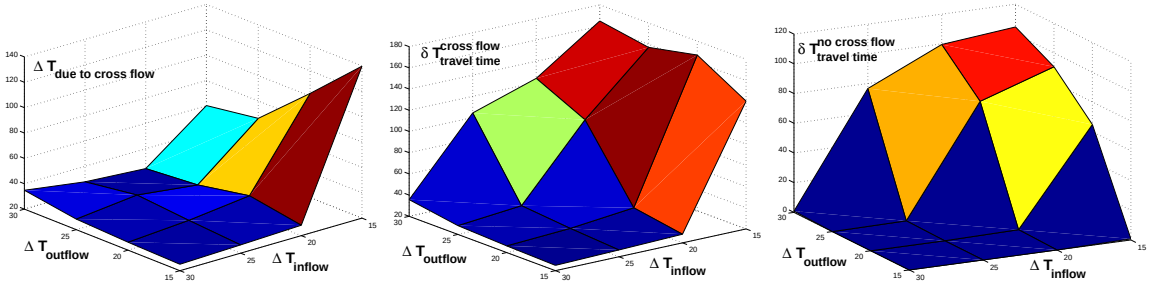

Figure 17: **Left:** Plot of $\Delta T_{\text{due to cross flow}}$ (in sec.) averaged over 10 runs. **Middle:** Plot of $\delta T^{\text{cross flow}}_{\text{travel time}}$ (in sec.) averaged over 10 runs. **Right:** Plot of $\delta T^{\text{no cross flow}}_{\text{travel time}}$ (in sec.) averaged over 10 runs.